Stochastic Privacy

Adish Singla
ETH Zurich
adish.singla@inf.ethz.ch

Eric Horvitz
Microsoft Research
horvitz@microsoft.com

Ece Kamar
Microsoft Research
eckamar@microsoft.com

Ryen White
Microsoft Research
ryen.white@microsoft.com

arXiv:1404.5454v1 [cs.AI] 22 Apr 2014

Abstract

Online services such as web search and e-commerce applications typically rely on the collection of data about users, including details of their activities on the web. Such personal data is used to enhance the quality of service via personalization of content and to maximize revenues via better targeting of advertisements and deeper engagement of users on sites. To date, service providers have largely followed the approach of either requiring or requesting consent for opting-in to share their data. Users may be willing to share private information in return for better quality of service or for incentives, or in return for assurances about the nature and extend of the logging of data. We introduce *stochastic privacy*, a new approach to privacy centering on a simple concept: A guarantee is provided to users about the upper-bound on the probability that their personal data will be used. Such a probability, which we refer to as *privacy risk*, can be assessed by users as a preference or communicated as a policy by a service provider. Service providers can work to personalize and to optimize revenues in accordance with preferences about privacy risk. We present procedures, proofs, and an overall system for maximizing the quality of services, while respecting bounds on allowable or communicated privacy risk. We demonstrate the methodology with a case study and evaluation of the procedures applied to web search personalization. We show how we can achieve near-optimal utility of accessing information with provable guarantees on the probability of sharing data.

Introduction

Online services such as web search, recommendation engines, social networks, and e-commerce applications typically rely on the collection of data about activities (*e.g.*, click logs, queries, and browsing information) and personal information (*e.g.*, location and demographics) of users. The availability of such data enables providers to personalize services to individuals and also to learn how to enhance the service for all users (*e.g.*, improved search results relevance). User data is also important to providers for optimizing revenues via better targeted advertising, extended user engagement and popularity, and even the selling of user data to third party companies. Permissions are typically obtained via broad consent agreements that request user permission to share their data through system dialogs, or via complex *Terms of Service*. Such notices are typically difficult to understand and ignored by more than 40 percent of users (Technet 2012). In other cases, a plethora of requests for information such as user location may be shown in system dialogs at run-time or installation time. Beyond the normal channels for sharing data, potential breaches of information are possible via attacks by malicious third parties and malware, and through surprising situations such as the AOL data release (Arrington 2006; Adar 2007) and de-anonymization of released Netflix logs (Narayanan and Shmatikov 2008). The charges by the Federal Trade Commission against Facebook (FTC 2011) and Google (FTC 2012) highlight increasing concerns by privacy advocates and government institutions about the large-scale recording of personal data.

Ideal approaches to privacy in online services would enable users to benefit from machine learning over data from populations of users, yet consider users' preferences as a top priority. Prior research in this realm has focused on designing privacy-preserving methodologies that can provide for control of a privacy-utility tradeoff (Adar 2007; Krause and Horvitz 2008). Research has also explored the feasibility of incorporating user preferences over what type of data can be logged (Xu et al. 2007; Cooper 2008; Olson, Grudin, and Horvitz 2005; Krause and Horvitz 2008).

We introduce a new approach to privacy that we refer to as *stochastic privacy*. Stochastic privacy centers on the simple idea of providing a guarantee to users about the maximum likelihood that their data will be accessed and used by a service provider. We refer to this measure as the assessed or communicated *privacy risk*, which may be increased in return for increases in the quality of service or other incentives. Very small probabilities of sharing data may be tolerated by individuals (just as lightning strikes are tolerated as a rare event), yet offer providers sufficient information to optimize over a large population of users. Stochastic privacy depends critically on harnessing inference and decision making to make choices about data collection within the constraints of a guaranteed privacy risk.

We explore procedures that can be employed by service providers when preferences about the sharing of data are represented as privacy risk. The goal is to maximize the

Copyright © 2014, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

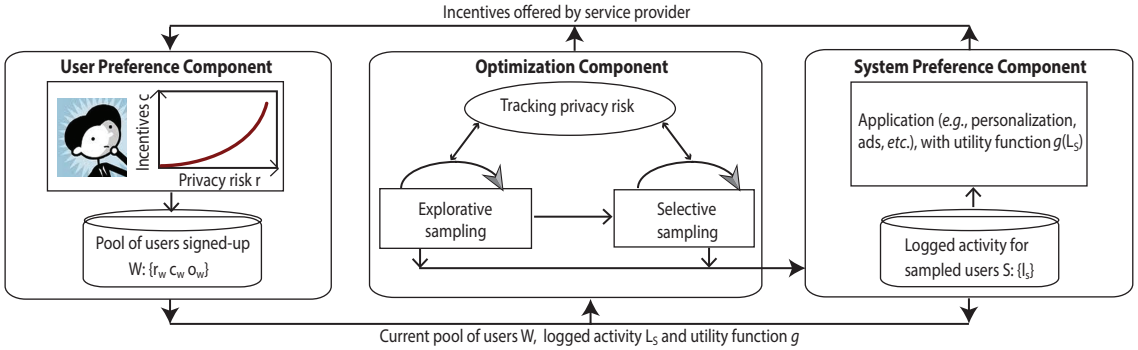
Incentives offered by service provider
User Preference Component
Incentives c
Privacy risk r
Pool of users signed-up
W: {r_w c_w o_w}
Optimization Component
Tracking privacy risk
Explorative sampling
Selective sampling
System Preference Component
Application (e.g., personalization, ads, etc.), with utility function g(L_S)
Logged activity for sampled users S: {l_s}
Current pool of users W, logged activity L_S and utility function g

Figure 1: Overview of stochastic privacy.

utility of service using data extracted from a population of users, while abiding by the agreement reached with users on privacy risk. We show that optimal selection of users under these constraints is NP-hard and thus intractable, given the massive size of the online systems. As a solution, we propose two procedures, RANDGREEDY and SPGREEDY, that combine greedy value of information analysis with obfuscation to offer mechanisms for tractable optimization, while satisfying stochastic privacy guarantees. We present performance bounds for the expected utility achievable by these procedures compared to the optimal solution. Our contributions can be summarized as follows:

- Introduction of stochastic privacy, an approach that represents preferences about the probability that data will be shared, and methods for trading off privacy risk, incentives, and quality of service.
- A tractable end-to-end system for implementing a version of stochastic privacy in online services.
- RANDGREEDY and SPGREEDY procedures for sampling users under the constraints of stochastic privacy, with theoretical guarantees on the acquired utility.
- Evaluation to demonstrate the effectiveness of the proposed procedures on a case study of user selection for personalization in web search.

Stochastic Privacy Overview

Figure 1 provides an overview of stochastic privacy in the context of a particular design of a system that implements the methodology. The design is composed of three main components: (i) a user preference component, (ii) a system preference component, and (iii) an optimization component for guiding the system's data collection. We now provide details about each of the components and then formally specify the optimization problem for *selective sampling* module.

User Preference Component

The user preference component interacts with users (*e.g.*, during signup) and establishes an agreement between a user and service provider on a tolerated probability that the user's data will be shared in return for better quality of service or incentives. Representing and capturing users' tolerated privacy risk allows users to move beyond the binary choice of yes or no on the sharing of data. The incentives offered to users can be personalized based on the metalevel information available for a user (*e.g.*, general location information inferred from a previously shared IP address) and can vary from guarantees of improved service (Krause and Horvitz 2010) to complementary software and entries in a lottery to win cash prizes (as done by the comScore service (Wikipedia-comScore 2006)).

Formally, let W be the population of users signed-up for a service. Each user $w \in W$ is represented with the tuple $\{r_w, c_w, o_w\}$, where o_w is the metadata information (*e.g.*, IP address) available for user w prior to selecting and logging finer-grained data about the user. r_w is the privacy risk assessed by the user, and c_w is the corresponding incentive provided in return for the user assuming the risk. The elements of this tuple can be updated through interactions between the system and the user. For simplicity of analysis, we shall assume that the pool W and user preferences are static.

System Preference Component

The goal of the service provider is to optimize the quality of service. For example, a provider may wish to personalize web search and to improve the targeting of advertising for maximization of revenue. The service provider may record the activities of a subset of users (*e.g.*, sets of queries issued, sites browsed, *etc.*) and use this data to provide better service globally or to a specific cohort of users. We model the private data of activity logs of user w by variable $l_w \in 2^L$, where L represents the web-scale space of activities (*e.g.*, set of queries issued, sites browsed, *etc.*) . However, l_w is observed by the system only after w is selected and the data from w is logged. We model the system's uncertain belief of l_w by a random variable Y_w, with l_w being its realization distributed according to conditional probability distribution $P(Y_w = l_w|o_w)$. In order to make an informed decision about user selection, the distribution $P(Y_w = l_w|o_w)$ is learned by the system using data available from the user and recorded logs of other users. We quantify the utility of application by logging activities L_S from selected users S through function $g : 2^L \rightarrow \mathbb{R}$, given by $g(\bigcup_{s \in S} l_s)$. The expected value of the utility that the system can expect to gain by selecting users S with observed attributes O_S is characterized by distribution P and utility function g as: $\tilde{g}(S) \equiv \mathbb{E}_{Y_S}\left[g(\bigcup_{s \in S} l_s)\right] = \sum_{L_S \in 2^L \times S} \left(P(Y_S = L_S|O_S) \cdot g(\bigcup_{s \in S} l_s)\right)$. However, the application itself may be using the logs L_S in a complex manner (such as training a ranker (Bennett et al. 2011)) and evaluating this on complex user metrics (Hassan and White 2013). Hence, the system

uses a surrogate utility function $f(S) \approx \tilde{g}(S)$ to capture the utility through a simple metric, for example, coverage of query-clicks obtained from the sampled users (Singla and White 2010) or reduction in uncertainty of click phenomena (Krause and Horvitz 2008).

In our model, we require that the set function f to be *non-negative*, *monotone* (i.e., whenever $A \subseteq A' \subseteq W$, it holds that $f(A) \leq f(A')$) and *submodular*. Submodularity is an intuitive notion of diminishing returns, stating that, for any sets $A \subseteq A' \subseteq W$, and any given user $a \notin A'$, it holds that $f(A \cup \{a\}) - f(A) \geq f(A' \cup \{a\}) - f(A')$. These conditions are general, satisfied by many realistic, as well as complex utility functions (Krause and Guestrin 2007), such as reduction in click entropy (Krause and Horvitz 2008). As a concrete example, consider the setting where attributes O represent geo-coordinates of the users and $D : O \times O \rightarrow \mathbb{R}$ computes the geographical distance between any two users. The goal of the service provider is to provide location-based personalization of web search. For such an application, click information from local users provides valuable signals for personalizing search (Bennett et al. 2011). The system's goal is to select a set of users S, and to leverage data from these users to enhance the service for the population. For search queries originating from any other user w, it uses the click data from the nearest user in S, given by $\arg\min_{s \in S} D(o_s, o_w)$. One approach for finding such a set S is solving the *k-medoid* problem which aims to minimize the sum of pairwise distances between selected set and the remaining population (Mirzasoleiman et al. 2013; Kaufman and Rousseeuw 2009). Concretely, this can be captured by the following submodular utility function:

$$f(S) = \frac{1}{|W|} \sum_{w \in W} \Big(\min_{x \in X} D(o_x, o_w) - \min_{s \in S \cup X} D(o_s, o_w) \Big) \tag{1}$$

Here, X is any one (or a set of) fixed reference location(s), for example, simply representing origin coordinates and is used ensure that function f is non-negative and monotone. Lemma 1 formally states the properties of this function.

Optimization Component

To make informed decisions about data access, the system computes the expected value of information (VOI) of logging the activities of a particular user, *i.e.*, the marginal utility that the application can expect by logging the activity of this user (Krause and Horvitz 2008). In the absence of sufficient information about user attributes, the VOI may be small, and hence needs to be learned from the data. The system can randomly sample a small set of users from the population that can be used to learn and improve the models of VOI computation (*explorative sampling* in Figure 1). For example, for optimizing the service for a user cohort speaking a specific language, the system may choose to collect logs from a subset of users to learn how languages spoken by users map to geography. If preferences about privacy risk were not being regarded, VOI can be used to select which users to log with a goal of maximizing the utility for the service provider (*selective sampling* in Figure 1). Given that the utility function of the system is submodular, a greedy selection rule makes near-optimal decisions about data access (Krause and Guestrin 2007). However, this simple approach could violate the privacy guarantees made with users. To act in accordance with the assessed privacy risk, we design selective sampling procedures that couple obfuscation with VOI analysis to select the set of users to provide data.

The system needs to ensure that both the explorative and selective sampling approaches respect the privacy guarantees made to users: the likelihood of sampling any user w throughout the execution of the system must be less than the privacy risk factor r_w. The system tracks the sampling risk (likelihood of sampling) that user w faces during phases of the execution of explorative sampling, denoted r_w^{ES}, and selective sampling, denoted r_w^{SS}. The privacy guarantee for a user is preserved as long as: $r_w - \big(1-(1-r_w^{ES})\cdot(1-r_w^{SS})\big) \geq 0$. This difference between the assessed risk and risk faced by a user can be viewed as the *sampling budget* of that user.

Optimization Problem for Selective Sampling

We now focus primarily on the selective sampling module and formally introduce the optimization problem. The goal is to design a sampling procedure M that abides by guarantees of stochastic privacy, yet optimizes the utility of the application in decisions about accessing user data. Given a budget constraint B, the goal is to select users S^M:

$$S^M = \arg\max_{S \subseteq W} f(S) \tag{2}$$
$$\text{subject to } \sum_{s \in S} c_s \leq B \text{ and } r_w - r_w^M \geq 0 \; \forall w \in W.$$

Here, r_w^M is the likelihood of selecting $w \in \mathcal{W}$ by procedure M and hence $r_w - r_w^M \geq 0$ captures the constraint of stochastic privacy guarantee for w. Note that we interchangeably write utility acquired by procedure as $f(M)$ to denote $f(S^M)$ where S^M is the set of users selected by running M. We shall now consider a simpler setting of constant privacy risk rate r for all users and unit cost per user (thus reducing the budget constraint to a simpler cardinal constraint, given by $|S| \leq B$). These assumptions lead to defining $B \leq W \cdot r$, as that is the maximum possible set size that can be sampled by any procedure for Problem 2.

Selective Sampling with Stochastic Privacy

We shall now propose desiderata of the selection procedures, discuss the hardness of the problem and review several different tractable approaches, as summarized in Table 1.

Desirable Properties of Sampling Procedures

The problem defined by Equation 2 requires solving an NP-hard discrete optimization problem, even when stochastic privacy constraint is removed. The algorithm for finding the optimal solution of this problem without the privacy constraint, referred as OPT, is intractable (Feige 1998). We address this intractability by exploiting the submodular structure of the utility function f and offer procedures providing provable near-optimal solutions in polynomial time. We aim at designing procedures that satisfy the following desirable properties: *(i)* provides competitive utility w.r.t. OPT with provable guarantees, *(ii)* preserves stochastic privacy guarantees, and *(iii)* runs in polynomial time.

Procedure	Competitive utility	Privacy guarantees	Polynomial runtime
OPT	✓	✗	✗ $\mathcal{O}(\lvert W\rvert^B)$
GREEDY	✓	✗	✓ $\mathcal{O}(B \cdot \lvert W\rvert)$
RANDOM	✗	✓	✓ $\mathcal{O}(B)$
RANDGREEDY	✓	✓	✓ $\mathcal{O}(B \cdot \lvert W\rvert \cdot r)$
SPGREEDY	✓	✓	✓ $\mathcal{O}(B \cdot \lvert W\rvert \cdot log(1/r))$

Table 1: Properties of different procedures. RANDGREEDY and SPGREEDY satisfy all the desirable properties.

Random Sampling: RANDOM

RANDOM simply samples the users at random, without any consideration of cost and utility. The likelihood of any user w to be selected by the algorithm is $r_w^{\text{RANDOM}} = B/W$ and hence privacy risk guarantees are trivially satisfied since $B \leq W \cdot r$ as defined in Problem 2). In general, RANDOM can perform arbitrarily poorly in terms of acquired utility, specifically for applications targeting particular user cohorts.

Greedy Selection: GREEDY

Next, we explore a greedy sampling strategy that maximizes the expected marginal utility at each iteration to guide the decision about selecting a next user to log. Formally, GREEDY starts with empty set $S = \emptyset$. At an iteration i, it greedily selects a user $s_i^* = \arg\max_{w \subseteq W \setminus S} f(S \cup w) - f(S)$ and adds it to the current selection of users $S = S \cup \{s_i^*\}$. It stops when $|S| = B$.

A fundamental result by Nemhauser, Wolsey, and Fisher (1978) states that the utility obtained by this greedy selection strategy is guaranteed to be at least $(1 - 1/e)$ $(= 0.63)$ times that obtained by OPT. This result is tight under reasonable complexity assumptions ($P \neq NP$) (Feige 1998). However, such a greedy selection clearly violates the stochastic privacy constraint in Problem 2—consider the user w^* with highest marginal value: $w^* = \arg\max_{w \subseteq W} f(w)$. The likelihood that this user will be selected by the algorithm $r_{w^*}^{\text{GREEDY}} = 1$, regardless of the requested privacy risk r_{w^*}.

Sampling and Greedy Selection: RANDGREEDY

We combine the ideas of RANDOM and GREEDY to design procedure RANDGREEDY which provides guarantees on stochastic privacy and competitive utility. RANDGREEDY is an iterative procedure that samples a small batch of users $\psi(s)$ at each iteration, then greedily selects $s^* \in \psi(s)$ and removes the entire set $\psi(s)$ for further consideration. By keeping the batch size $\psi(s) \leq W \cdot r/B$, the procedure ensures that the privacy guarantees are satisfied. As our user pool W is static, to reduce complexity, we consider a simpler version of RANDGREEDY that defers the greedy selection. Formally, this is equivalent to first sampling the users from W at rate r to create a subset $\widetilde{W}$ such that $|\widetilde{W}| = |W| \cdot r$, and then, running the GREEDY algorithm on $\widetilde{W}$ to greedily select a set of users of size B.

The initial random sampling ensures a guarantee on the privacy risk for users during the execution of the procedure. In fact, for any user $w \in W$, the likelihood of w being sampled and included in subset $\widetilde{W}$ is $r_w^{\text{RANDGREEDY}} \leq r$. We further analyze the utility obtained by this procedure in the next section and show that, under reasonable assumptions, the approach can provide competitive utility compared to OPT.

Greedy Selection with Obfuscation: SPGREEDY

SPGREEDY uses an inverse approach of mixing RANDOM and GREEDY: it does greedy selection, followed by obfuscation, as illustrated in Procedure 1. It assumes an underlying distance metric $D : W \times W \to \mathbb{R}$ which captures the notion of distance or dissimilarity among users. As in GREEDY, it operates in iterations and selects the element s^* with maximum marginal utility at each iteration. However, to ensure stochastic privacy, it obfuscates s^* with similar users using distance metric D to create a set $\psi(s^*)$ of size $1/r$, then samples one user randomly from $\psi(s^*)$ and removes the entire set $\psi(s^*)$ for further consideration.

The guarantees on privacy risk hold by the following arguments: During the execution of the algorithm, any user w becomes a possible candidate of being selected if the user is part of $\psi(s^*)$ in some iteration (*e.g.*, iteration i). Given that $|\psi(s^*)| \geq 1/r$ and algorithm randomly sample $v \in \psi(s^*)$, the likelihood of w being selected in iteration i is at most r. The fact that set $\psi(s^*)$ is removed from available pool $\widetilde{W}$ at the end of the iteration ensures that w can become a possible candidate of selection only once.

Procedure 1: SPGREEDY

Input: *users* W; *cardinality constraint* B; *privacy risk* r; *distance metric* $D : W \times W \to \mathbb{R}$;
Initialize:
- **Outputs**: *selected users* $S \leftarrow \emptyset$;
- **Variables**: *remaining users* $W' \leftarrow W$;

begin
  while $|S| \leq B$ **do**
    $s^* \leftarrow \arg\max_{w \in W'} f(S \cup w) - f(S)$;
    Set $\psi(s^*) \leftarrow s^*$;
    while $|\psi(s^*)| < 1/r$ **do**
      $v \leftarrow \arg\min_{w \in W' \setminus \psi(s^*)} D(w, s^*)$;
      $\psi(s^*) \leftarrow \psi(s^*) \cup \{v\}$;
    Randomly select $\widetilde{s^*} \in \psi(s^*)$;
    $S \leftarrow S \cup \{\widetilde{s^*}\}$;
    $W' \leftarrow W' \setminus \psi(s^*)$;
Output: S

Performance Analysis

We now analyze the performance of the proposed procedures in terms of the utility acquired compared to that of the OPT as baseline. We first analyze the problem in a general setting and then under a set of practical assumptions on the structure of underlying utility function f and population of users W. The proofs of all the results are available at [1].

[1] Available anonymously at:
`http://tinyurl.com/aaai-stocpriv-longer`

General Case

In the general setting, we show that one cannot do better than $r \cdot f(OPT)$ in the worst case. Consider a population of users W where only one user w^* has utility value of 1, and rest of the users $W \setminus w^*$ have utility of 0. The OPT gets a utility of 1 by selecting $S^{\text{OPT}} = \{w^*\}$. Consider any procedure M that has to respect the guarantees on privacy risk. If the privacy rate of w^* is r, then M can select w^* with only a maximum probability of r. Hence, the maximum expected utility that any procedure M for Problem 2 can achieve is r.

On a positive note, a trivial algorithm can always achieve a utility of $(1 - 1/e) \cdot r \cdot f(OPT)$ in expectation. This result can be achieved by running GREEDY to select a set S^{GREEDY} and then choosing the final solution to be S^{GREEDY} with probability r, and else output an empty set. Theorem 1 formally states these results for the general problem setting.

Theorem 1. *Consider the Problem 2 of optimizing a submodular function f under cardinality constraint B and privacy risk rate r. For any distribution of marginal utilities of population W, a trivial procedure can achieve expected utility of at least $(1 - 1/e) \cdot r \cdot f(OPT)$. In contrast, there exists an underlying distribution for which no procedure can have expected utility of more than $r \cdot f(OPT)$.*

Smoothness and Diversification Assumptions

In practice, we can hope to do much better than the worst-case results described in Theorem 1 by exploiting the underlying structures of users attributes and utility function. We start with the assumption that there exists a distance metric $D : W \times W \to \mathbb{R}$ which captures the notion of distance or dissimilarity among users. For any given $w \in W$, let us define its α-neighborhood to be the set of users within a distance α from w (i.e., α-close to w): $N_\alpha(w) = \{v : D(v, w) \leq \alpha\}$. We assume that population of users is large and that the number of users in the $N_\alpha(w)$ is large. We formally capture these requirements in Theorems 2,3.

Firstly, we consider utility functions that change gracefully with changes in inputs, similar to the notion of λ-*Lipschitz* set functions used in Mirzasoleiman et al. (2013). We formalize the notion of smoothness in the utility function f w.r.t metric D as follows:

Definition 1. *For any given set of users S, let us consider a set $\widetilde{S}_\alpha$ obtained by replacing every $s \in S$ with any $w \in N_\alpha(s)$. Then, $|f(S) - f(\widetilde{S}_\alpha)| \leq \lambda_f \cdot \alpha \cdot |S|$, where parameter λ_f captures the notion of smoothness of function f.*

Secondly, we consider utility functions that favor diversity or dissimilarity of users in the subset selection w.r.t the distance metric D. We formalize this notion of diversification in the utility function as follows:

Definition 2. *Let us consider any given set of users $S \subseteq W$ and a user $w \in W$. Let $\alpha = \min_{s \in S} D(s, w)$. Then, $f(S \cup w) - f(S) \leq \Upsilon_f \cdot \alpha$, where parameter Υ_f captures the notion of diversification of function f.*

The utility function f introduced in Equation 1 satisfies both the above assumptions as formally stated below.

Lemma 1. *Consider the utility function f in Equation 1. f is submodular, and satisfies the properties of smoothness and diversification, i.e. has bounded λ_f and Υ_f.*

We note that for the functions with unbounded λ and Υ (*i.e.*, $\lambda_f \to \infty$ and $\Upsilon_f \to \infty$), it would lead to the general problem settings (equivalent to no assumptions) and hence results of Theorem 1 apply.

Performance Bounds

Under the assumption of smoothness (*i.e.*, bounded λ_f), we can show the following bound on utility of RANDGREEDY:

Theorem 2. *Consider the Problem 2 for function f with bounded λ_f. Let S^{OPT} be the set returned by* OPT *for Problem 2 after relaxing privacy constraints. For a desired $\epsilon < 1$, let $\alpha_{rg} = \arg\min_\alpha \{\alpha : |N_\alpha(s)| \geq 1/r \cdot \log(B/\epsilon)\ \forall s \in S^{\text{OPT}}$, where $N_\alpha(s_i) \cap N_\alpha(s_j) = \emptyset\ \forall s_i, s_j \in S^{\text{OPT}}\}$. Then, with probability at least $(1 - \epsilon)$,*

$$\mathbb{E}[f(\text{RANDGREEDY})] \geq (1 - 1/e) \cdot \big(f(\text{OPT}) - \alpha_{rg} \cdot \lambda_f \cdot B\big).$$

Under the assumption of smoothness and diversification (*i.e.*, bounded λ_f and Υ_f), we can show the following bound on utility of SPGREEDY:

Theorem 3. *Consider the Problem 2 for function f with bounded λ_f and Υ_f. Let S^{GREEDY} be the set returned by* GREEDY *for Problem 2. Let $\alpha_{spg} = \arg\min_\alpha \{\alpha : |N_\alpha(s)| \geq 1/r\ \forall s \in S^{\text{GREEDY}}\}$. Then,*

$$\mathbb{E}[f(\text{SPGREEDY})] \geq (1 - 1/e) \cdot f(\text{OPT}) - (2 \cdot \lambda_f + \Upsilon_f) \cdot \alpha_{spg} \cdot B.$$

Intuitively, these results imply that both RANDGREEDY and SPGREEDY achieve competitive utility w.r.t OPT, and the performance degrades smoothly as the privacy risk r is decreased or the bounds on smoothness and diversification increase.

Experimental Evaluation

We shall now report on experiments we performed to build insights about the performance of the stochastic privacy procedures on a case study of the selective collection of user data in support of web search personalization.

Benchmarks and Metrics

We compare the performance of the RANDGREEDY and SPGREEDY procedures against the baselines of RANDOM and GREEDY. While RANDOM provides a trivial lower benchmark for any procedure, GREEDY is a natural upper bound on the utility, given that the OPT itself is intractable. To analyze the robustness of our procedures, we then vary the level of privacy risk r. We further carried out experiments to understand the loss incurred from obfuscation phase during the execution of SPGREEDY.

Experimental Setup

We considered the application of providing location-based personalization for queries issued for the business domain (*e.g.*, real-estate, financial services, *etc.*). The goal is to select a set of users S who are expert web search users in this domain. We seek to leverage the click data from these users to improve the relevance of search results shown to those searching for local businesses. The experiments are based on using a surrogate utility function as introduced in Equation 1. As we are interested in specific domain of business-related queries, we modify the utility function in Equation 1

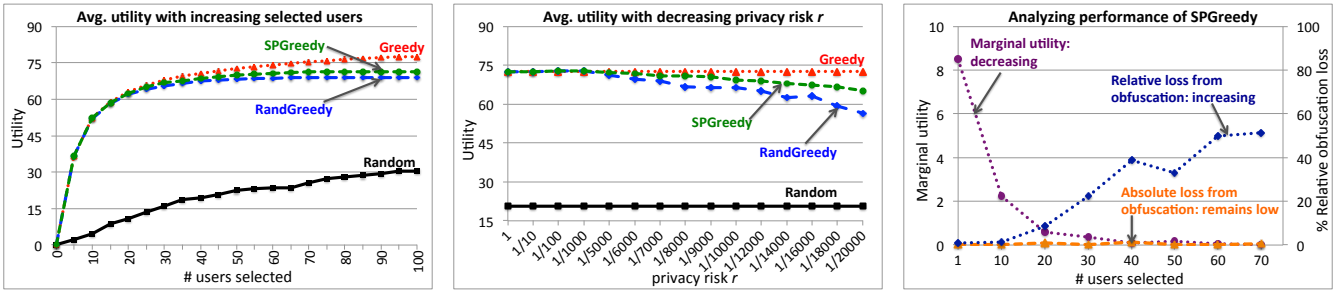

(a) Vary budget B (b) Vary privacy risk r (c) Analyze SPGREEDY

Figure 2: In Fig. 2(a), for a fixed $r = 1/10000$, budget B or number of users selected in increased, showing the competitiveness of our procedures w.r.t GREEDY. In Fig. 2(b), a fixed $B = 50$ is used, and level of privacy risk r is reduced. The results demonstrate that the performance of RANDGREEDY and SPGREEDY degrades smoothly. Fig. 2(c) analyze the execution of the procedure SPGREEDY and illustrates that the loss incurred in marginal utility at every step via obfuscation is very low.

by restricting S to users who are experts in the domain, as further described below. The acquired utility can be interpreted as the average reduction in the distance for any user w in the population to the nearest expert $s \in S$.

The primary source of data for this study is obtained from interaction logs on a major web search engine. We considered a fraction of users who issued at least one query in month of October 2013, restricted to queries coming from IP addresses located within ten neighboring states in the western region of the United States. This resulted in a pool W of seven million users. We considered a setting where system has access to metadata information of geo-coordinates of the users, as well as a probe of the last 20 search-result clicks for each user, which together constitutes the observed attributes of user denoted as o_w. Each of these clicks are then classified into a topical hierarchy from a popular web directory, the Open Directory Project (ODP) (dmoz.org), using automated techniques (Bennett, Svore, and Dumais 2010). With a similar objective to White, Dumais, and Teevan (2009), the system then uses this classification to identify users who are expert in the business domain. We used a simple rule of classifying a user as an expert if at least one click was issued in the domain of interest. With this, the system marks a set of users $W' \subseteq W$ as experts, and the set S in Equation 1 is restricted to W'. We note that the specific thresholds or variable choices do not affect the overall results below.

Results

We now discuss the findings from our experiments.

Varying the budget B: In our first set of experiments, we vary the budget B, or equivalently the number of users selected, and measured the utility acquired by different procedures. The privacy risk rate is set fixed to $r = 1/10000$. Figure 2(a) illustrates that both RANDGREEDY and SPGREEDY are competitive w.r.t GREEDY and clearly outperform naive baseline of RANDOM.

Varying the privacy risk r: We then vary the level of privacy risk, for a fixed budget $B = 50$, to measure the robustness of the RANDGREEDY and SPGREEDY. The results in Figure 2(b) demonstrate that the performance of RANDGREEDY and SPGREEDY degrades smoothly, as per the performance analysis in Theorems 2,3.

Analyzing performance of SPGREEDY: Lastly, we perform experiments to understand the execution of SPGREEDY and the loss incurred from the obfuscation step. SPGREEDY removes $1/r$ users from pool at every iteration. As a result, for small privacy risk r, the relative loss from obfuscation (*i.e.*, relative % difference in marginal utility acquired by a user chosen by greedy selection, compared to one finally picked after obfuscation) could possibly increase over the execution of procedure as illustrated in Figure 2(a), using a moving average of window size 10. However, the diminishing returns property of the utility ensures that SPGREEDY incurs very low absolute loss in marginal utility from obfuscation at every step.

Discussion

We introduced *stochastic privacy*, a new approach to privacy that centers on service providers abiding by guarantees about not exceeding a specified likelihood of logging data, and maximizing information collection in accordance with these guarantees. We presented procedures and an overall system design for maximizing the quality of services while respecting privacy risks agreed with populations of users.

Directions for this research include the assessments of user preferences about the probability of sharing data, including how users trade increases in privacy risk for enhanced service and monetary incentives. Directions also include exploring the rich space of designs for interactive and longer-term controls and settings of a tolerated risk of sharing data. Opportunities include policies and analyses based on the sharing of data as a privacy risk *rate* over time. As an example, systems might one day consider decisions about logging one or more search sessions of users where privacy risk is assessed in terms of the risk of sharing search sessions over time. In another research direction, designs can include models where users are notified when they are selected to share data and are provided with a special reward and option of declining to share at that time. Iterative analyses can be developed where subsets of users are actively engaged with the option to assume higher levels of privacy risk or to simply provide additional information in return for special incentives. Inferences about the likely preferences on privacy risk and about incentives for subpopulations could be folded into the selection procedures.

References

[2007] Adar, E. 2007. User 4xxxxx9: Anonymizing query logs. In *Workshop on Query Log Analysis at WWW'07*.

[2006] Arrington, M. 2006. Aol proudly releases massive amounts of private data. http://techcrunch.com/2006/08/06/aol-proudly-releases-massive-amounts-of-user-search-data/.

[2011] Bennett, P. N.; Radlinski, F.; White, R. W.; and Yilmaz, E. 2011. Inferring and using location metadata to personalize web search. In *Proc. of SIGIR*, 135–144.

[2010] Bennett, P. N.; Svore, K.; and Dumais, S. T. 2010. Classification-enhanced ranking. In *Proc. of WWW*, 111–120.

[2008] Cooper, A. 2008. A survey of query log privacy-enhancing techniques from a policy perspective. *ACM Trans. Web* 2(4):19:1–19:27.

[1998] Feige, U. 1998. A threshold of ln n for approximating set cover. *Journal of the ACM* 45:314–318.

[2011] FTC. 2011. FTC charges against Facebook. http://www.ftc.gov/opa/2011/11/privacysettlement.shtm.

[2012] FTC. 2012. FTC charges against Google. http://www.ftc.gov/opa/2012/08/google.shtm.

[2013] Hassan, A., and White, R. W. 2013. Personalized models of search satisfaction. In *Proc. of CIKM*, 2009–2018.

[2009] Kaufman, L., and Rousseeuw, P. J. 2009. *Finding groups in data: an introduction to cluster analysis*, volume 344. Wiley.com.

[2005] Krause, A., and Guestrin, C. 2005. A note on the budgeted maximization on submodular functions. Technical Report CMU-CALD-05-103, Carnegie Mellon University.

[2007] Krause, A., and Guestrin, C. 2007. Near-optimal observation selection using submodular functions. In *Proc. of AAAI, Nectar track*.

[2008] Krause, A., and Horvitz, E. 2008. A utility-theoretic approach to privacy and personalization. In *Proc. of AAAI*.

[2010] Krause, A., and Horvitz, E. 2010. A utility-theoretic approach to privacy in online services. *Journal of Artificial Intelligence Research (JAIR)* 39:633–662.

[2013] Mirzasoleiman, B.; Karbasi, A.; Sarkar, R.; and Krause, A. 2013. Distributed submodular maximization: Identifying representative elements in massive data. In *Proc. of NIPS*.

[2008] Narayanan, A., and Shmatikov, V. 2008. Robust de-anonymization of large sparse datasets. In *Proc. of the IEEE Symposium on Security and Privacy*, 111–125.

[1978] Nemhauser, G.; Wolsey, L.; and Fisher, M. 1978. An analysis of the approximations for maximizing submodular set functions. *Math. Prog.* 14:265–294.

[2005] Olson, J.; Grudin, J.; and Horvitz, E. 2005. A study of preferences for sharing and privacy. In *Proc. of CHI*.

[2010] Singla, A., and White, R. W. 2010. Sampling high-quality clicks from noisy click data. In *Proc. of WWW*, 1187–1188.

[2012] Technet. 2012. Privacy and technology in balance. http://blogs.technet.com/b/microsoft_on_the_issues/archive/2012/10/26/privacy-and-technology-in-balance.aspx.

[2009] White, R. W.; Dumais, S. T.; and Teevan, J. 2009. Characterizing the influence of domain expertise on web search behavior. In *Proc. of WSDM*, 132–141.

[2006] Wikipedia-comScore. 2006. ComScore-#Data_collection_and_reporting. http://en.wikipedia.org/wiki/ComScore#Data_collection_and_reporting.

[2007] Xu, Y.; Wang, K.; Zhang, B.; and Chen, Z. 2007. Privacy-enhancing personalized web search. In *Proc. of WWW*, 591–600. ACM.

Proof of Lemma 1

We prove Lemma 1 by proving three other Lemmas 2 3 4 that are not in the main paper. In Lemma 2, by using the decomposable property of the function f from Equation 1, we prove that the function f is non-negative, monotonous (non-decreasing) and submodular. Then, we show that the function satisfies the properties of smoothness (in Lemma 3) and diversification (in Lemma 4) by showing an upper bound on the values of the parameters λ_f and Υ_f.

Lemma 2. *Utility function f in Equation 1 is non-negative, monotone (non-decreasing) and submodular.*

Proof. We begin by noting that f is *decomposable*, *i.e.*, it can be written as a sum of simpler functions f_w as:

$$f(S) = \sum_{w \in W} f_w(S) \tag{3}$$

where $f_w(S)$ is given by:

$$f_w(S) = \frac{1}{|W|}\Big(\min_{x \in X} D(o_x, o_w) - \min_{s \in S \cup X} D(o_s, o_w)\Big) \tag{4}$$

Next, we prove that each of these functions f_w is non-negative, non-decreasing and submodular. To prove that the function is non-decreasing, consider any two sets $S \subseteq S' \subseteq W$. Then,

$$\begin{aligned} f_w(S') - f_w(S) &= \frac{1}{|W|}\Big(\min_{s \in S \cup X} D(o_s, o_w) - \min_{s \in S' \cup X} D(o_s, o_w)\Big) \\ &\geq 0 \end{aligned} \tag{5}$$

In step 5, the inequality holds as the distance to the nearest user for w in S' cannot be more than that in S, hence proving that f_w is non-decreasing. Also, it is easy to see that $f_w(\emptyset) = 0$, which along with the non-decreasing property, ensures that the function f_w is non-negative.

To prove that the function is submodular, consider any two sets $S \subseteq S' \subseteq W$, and any given user $v \in W \setminus S'$. When $f_w(S' \cup \{v\}) - f_w(S') = 0$, submodularity holds trivially as we have $f_w(S \cup \{v\}) - f_w(S) \geq 0$ using non-decreasing property. Let us consider the case when $f_w(S' \cup \{v\}) - f_w(S') > 0$, *i.e.*, v is assigned as the nearest user to w from the set $S' \cup \{v\}$, given by $v = \min_{s \in S' \cup \{v\} \cup X} D(o_s, o_w)$. In this case, it would also be the case that v is the nearest user to w from the set $S \cup \{v\}$. Then, we can write down the difference of marginal gains as follows:

$$\begin{aligned} &\Big(f_w(S' \cup \{v\}) - f_w(S')\Big) - \Big(f_w(S \cup \{v\}) - f_w(S)\Big) \\ &\qquad = \left(\frac{1}{|W|}\Big(D(o_v, o_w) - \min_{s \in S' \cup X} D(o_s, o_w)\Big)\right) - \left(\frac{1}{|W|}\Big(D(o_v, o_w) - \min_{s \in S \cup X} D(o_s, o_w)\Big)\right) \\ &\qquad = \frac{1}{|W|}\Big(\min_{s \in S \cup X} D(o_s, o_w) - \min_{s \in S' \cup X} D(o_s, o_w)\Big) \\ &\qquad \leq 0 \end{aligned} \tag{6}$$

In step 6, the inequality holds as the function is non-decreasing, thus showing that the marginal gains diminish and hence proving the submodularity of the function f_w.

By using the fact that these properties are preserved under linear combination with non-negative weights (all equal to 1 from Equation 3), f is non-negative, non-decreasing and submodular. □

Lemma 3. *Utility function f in Equation 1 satisfies the properties of smoothness, i.e. has bounded λ_f.*

Proof. For any given set of users S, let us consider a set $\widetilde{S}_\alpha$ obtained by replacing every $s \in S$ with any $w \in N_\alpha(s)$. The goal is to show that $|f(S) - f(\widetilde{S}_\alpha)| \leq \lambda_f \cdot \alpha \cdot |S|$ always holds for a fixed and bounded λ_f.

Let us again use the simpler functions f_w from decomposition of f in Equation 3 and consider the difference $|f_w(S) - f_w(\widetilde{S}_\alpha)|$. Then,

$$|f_w(S) - f_w(\widetilde{S}_\alpha)| \leq \frac{\alpha}{|W|} \tag{7}$$

In step 7, the inequality holds as the deviation in distance to the nearest user for w in $\widetilde{S}_\alpha$ cannot be more than α. Using this result, we have

$$|f(S) - f(\widetilde{S}_\alpha)| = |\sum_{w \in W} f_w(S) - \sum_{w \in W} f(\widetilde{S}_\alpha)|$$

$$\leq \sum_{w \in W} |f_w(S) - f_w(\widetilde{S}_\alpha)|$$
$$\leq \sum_{w \in W} \frac{\alpha}{|W|} \tag{8}$$
$$= \alpha \leq \alpha \cdot |S| \tag{9}$$

The inequality in step 8 holds by using the result of step 7 and inequality in step 9 holds trivially as $|S| \geq 1$. Hence, the smoothness parameter of the function λ_f is bounded by 1. □

Lemma 4. *Utility function f in Equation 1 satisfies the properties of diversification, i.e. has bounded Υ_f.*

Proof. For any given set of users S and any new user $v \in W \setminus S$, let us define $\alpha = \min_{s \in S} D(s, v)$. The goal is to show that $f(S \cup v) - f(S) \leq \Upsilon_f \cdot \alpha$ always holds for a fixed and bounded Υ_f.

Again, let us consider the function f_w and consider the marginal of adding v to S, given by $f_w(S \cup v) - f_w(S)$. When v is not the nearest user to w in the set $S \cup \{v\}$, we have $f(S \cup \{v\}) - f(S) = 0$. Let's consider the case where $f_w(S \cup v) - f(S) > 0$, *i.e.*, v is assigned as the nearest user to w from the set $S \cup \{v\}$, given by $v = \min_{s \in S \cup \{v\} \cup X} D(o_s, o_w)\Big)$. Let us denote the nearest user assigned to w before adding v to the set by v'. Then, we have:

$$f_w(S \cup v) - f_w(S) = \frac{1}{|W|}\big(D(s, v') - D(s, v)\big)$$
$$\leq \frac{D(v, v')}{|W|} \tag{10}$$
$$\leq \frac{\alpha}{|W|} \tag{11}$$

Step 10 uses the triangular inequality of the underlying metric space. In step 11, the inequality holds by the definition of α. Then, we have

$$f(S \cup v) - f(S) = \sum_{w \in W} \big(f_w(S \cup v) - \sum_{w \in W} f(S)\big)$$
$$\leq \sum_{w \in W} \frac{\alpha}{|W|} \tag{12}$$
$$= \alpha$$

The inequality in step 12 holds by using the result of step 11. Hence, the diversification parameter of the function Υ_f is bounded by 1. □

Proof of Lemma 1. The proof directly follows from the results in Lemmas 2 3 4. □

Proof of Theorem 2

Proof of Theorem 2. Let S^{OPT} be the set returned by OPT for Problem 2 without the privacy constraints. By the hypothesis of the theorem, for each of the element $s \in S^{\text{OPT}}$, the α_{rg} neighborhood of s contains a set of at least $1/r \cdot \log(B/\epsilon)$ users. Furthermore, by hypothesis, these sets of size at least $1/r \cdot \log(B/\epsilon)$ can be constructed to be mutually disjoint for every $s, s' \in S^{\text{OPT}}$, let us denote these mutually disjoint sets by $\widetilde{N}_{\alpha_{rg}}(s)$. Formally, this means that for $s \in S^{\text{OPT}}$, we have $|\widetilde{N}_{\alpha_{rg}}(s)| \geq 1/r \cdot \log(B/\epsilon)$ and for any pairs of $s, s' \in S^{\text{OPT}}$, we have $\widetilde{N}_{\alpha_{rg}}(s) \cap \widetilde{N}_{\alpha_{rg}}(s) = \emptyset$.

Recall that the simpler version of RANDGREEDY first samples the users from W at rate r to create a subset $\widetilde{W}$ such that $|\widetilde{W}| = |W| \cdot r$. We first show that sampling at a rate r by RANDGREEDY ensures that with high probability (given by $1 - \epsilon$), at least one user is sampled from $\widetilde{N}_{\alpha_{rg}}(s)$ for each of the $s \in S^{\text{OPT}}$. Consider the process of sampling for s and $\widetilde{N}_{\alpha_{rg}}(s)$. Each of the users in $\widetilde{N}_{\alpha_{rg}}(s)$ has probability of being sampled given by r. Hence, the probability that none of the users in $\widetilde{N}_{\alpha_{rg}}(s)$ are included in $\widetilde{W}$ for a given s is given by:

$$P\big(\widetilde{N}_{\alpha_{rg}}(s) \cap \widetilde{W} = \emptyset\big) = (1 - r)^{1/r \cdot \log(B/\epsilon)}$$
$$\leq e^{-\log(B/\epsilon)}$$
$$= \epsilon/B$$

By using union bound, the probability that none of the users in $\widetilde{N}_{\alpha_{rg}}(s)$ gets included in $\widetilde{W}$ for any $s \in S^{\text{OPT}}$ is bounded by ϵ (given by $B \cdot \epsilon/B$). Hence, with probability at least $1-\epsilon$, the sampled set $\widetilde{W}$ contains at least one user from $\widetilde{N}_{\alpha_{rg}}(s)$ for every $s \in S^{\text{OPT}}$.

This is equivalent to saying that, with probability at least $1-\epsilon$, the $\widetilde{W}$ contains a set $\widetilde{S}^{\text{OPT}}_{\alpha_{rg}}$ that can be obtained by replacing every $s \in S^{\text{OPT}}$ with some $w \in N_{\alpha_{rg}}(s)$, and hence $f(\widetilde{S}^{\text{OPT}}_{\alpha_{rg}}) \geq f(S^{\text{OPT}}) - \alpha_{rg} \cdot \lambda_f \cdot B$ (by using the definition of smoothness property). And, running the GREEDY on $\widetilde{W}$ ensures that the utility obtained is at least $(1-1/e) \cdot f(\widetilde{S}^{\text{OPT}}_{\alpha_{rg}})$. Hence, with probability at least $(1-\epsilon)$,

$$\begin{aligned}\mathbb{E}[f(\text{RANDGREEDY})] &\geq (1-1/e) \cdot f(\widetilde{S}^{\text{OPT}}_{\alpha_{rg}}) \\ &\geq (1-1/e) \cdot \big(f(\text{OPT}) - \alpha_{rg} \cdot \lambda_f \cdot B\big)\end{aligned}$$

□

Proof of Theorem 3

Proof of Theorem 3. Let S^{GREEDY} be the set returned by GREEDY for Problem 2 without the privacy constraints. By the hypothesis of the theorem, for each of the element $s \in S^{\text{GREEDY}}$, the α_{spg} neighborhood of s contains a set of at least $1/r$ users. The loss of utility for the procedure SPGREEDY compared w.r.t to GREEDY at iteration i can be attributed to two following reasons: (1) obfuscation of s^*_i with set $\psi(s^*_i)$ to select $\widetilde{s^*_i}$, where the size of $\psi(s^*_i)$ is $1/r$, and (2) removal of the entire set $\psi(s^*_i)$ for further consideration. We analyze these two factors separately to get the desired bounds on the utility of SPGREEDY.

We being by stating a more general result on the approximation guarantees of GREEDY from (Krause and Guestrin 2005) when the submodular objective function can only be evaluated approximately within an absolute error of ϵ. Results from (Krause and Guestrin 2005) states that the utility obtained by this noisy greedy selection is guaranteed to be at least $\big((1-1/e) \cdot \text{OPT} - 2 \cdot \epsilon \cdot B\big)$, where B is the budget.

Now, consider an alternate procedure that operates similar to SPGREEDY, by obfuscating s^*_i with set $\psi(s^*_i)$ to pick $\widetilde{s^*_i}$ at each iteration i. However, this alternate procedure does not eliminate the entire set of users $\psi(s^*_i)$ from the pool, but only removes $\widetilde{s^*_i}$. Instead, it tags the users of $\psi(s^*_i) \setminus \{\widetilde{s^*_i}\}$ as $<invalid, i>$, *i.e.* these users are marked as *invalid* and are tagged with the iteration i at which they became invalid (in case a user was already marked as invalid, the iteration tag is not updated). Let us denote this alternate procedure by $\overline{\text{SPGREEDY}}$. This can alternatively be viewed as similar to GREEDY, though it can pick the user at every iteration only approximately, because of the noise added by obfuscation. We now bound the absolute value of this approximation error at every iteration. As s^*_i is obfuscated with a set of users of size $1/r$ nearest to s^*_i from the hypothesis of the theorem, we are certain that set $\psi(s^*_i)$ is contained within a radius of α_{spg} neighborhood. Now, from the smoothness assumptions, the maximum absolute error that could be introduced by the obfuscation compared to greedy selection (*i.e.* the difference in marginal utilities of s^*_i and $\widetilde{s^*_i}$) at a given iteration i is bounded by $\lambda_f \cdot \alpha_{spg}$. Hence, the utility obtained by $\overline{\text{SPGREEDY}}$ can be lower-bounded as:

$$f(\overline{\text{SPGREEDY}}) \geq (1-1/e) \cdot \text{OPT} - 2 \cdot \lambda_f \cdot \alpha_{spg} \cdot B \tag{13}$$

Next, we consider the loss associated with the removal of entire set $\psi(s^*_i)$ at iteration i. Let us consider the execution of $\overline{\text{SPGREEDY}}$ and let $l+1$ be the first iteration when the obfuscation set $\psi(s^*_{l+1})$ created by the procedure contains at least one element marked as *invalid*, with the associated iteration of invalidity as k. Note that when $l+1 > B$, there is no loss associated with this step of removing $\psi(s^*_i)$ and hence we only consider the case when $l+1 \leq B$. As the users are embedded in euclidean space, this means that the α_{spg} centered around s^*_{l+1} and s^*_k overlaps and hence $D(s^*_{l+1}, s^*_k) \leq 2 \cdot \alpha_{spg}$. From the diversification assumption, this means that the marginal utility of s^*_{l+1} cannot be more than $2 \cdot \upsilon_f \cdot \alpha_{spg}$. And, furthermore, the submodularity ensures that for all $j > l+1$, the marginal utility of users selected can only be lesser than the marginal utility of s^*_{l+1}.

Let us consider a truncated version of $\overline{\text{SPGREEDY}}$ that stops after l steps, denoted by $\overline{\text{SPGREEDY}}_V$, where V denotes the fact that this procedure is always valid as it never touches *invalid* marked users. The utility of the truncated version can be lower-bounded as follows:

$$\begin{aligned}f(\overline{\text{SPGREEDY}}_V) &\geq f(\overline{\text{SPGREEDY}}) - (B-l) \cdot (2 \cdot \upsilon_f \cdot \alpha_{spg}) \\ &\geq f(\overline{\text{SPGREEDY}}) - 2 \cdot \upsilon_f \cdot \alpha_{spg} \cdot B \\ &\geq (1-1/e) \cdot \text{OPT} - (2 \cdot \lambda_f + 2 \cdot \upsilon_f) \cdot \alpha_{spg} \cdot B \end{aligned} \tag{14}$$

The step 14 follows by using the result in step 13. For the first l iterations, the execution of the mechanism SPGREEDY is exactly same as $\overline{\text{SPGREEDY}}_V$. Hence, SPGREEDY acquires utility at least that acquired by $\overline{\text{SPGREEDY}}_V$, which completes the proof. □